Asymptotics of Fingerprinting and Group Testing: Tight Bounds from Channel Capacities

Thijs Laarhoven

arXiv:1404.2576v1 [cs.IT] 9 Apr 2014

***Abstract*—In this work we consider the large-coalition asymptotics of various fingerprinting and group testing games, and derive explicit expressions for the capacities for each of these models. We do this both for simple decoders (fast but suboptimal) and for joint decoders (slow but optimal).**

For fingerprinting, we show that if the pirate strategy is known, the capacity often decreases linearly with the number of colluders, instead of quadratically as in the uninformed fingerprinting game. For many attacks the joint capacity is further shown to be strictly higher than the simple capacity.

For group testing, we improve upon known results about the joint capacities, and derive new explicit asymptotics for the simple capacities. These show that existing simple group testing algorithms are suboptimal, and that simple decoders cannot asymptotically be as efficient as joint decoders. For the traditional group testing model, we show that the gap between the simple and joint capacities is a factor $\log_2(e) \approx 1.44$ for large numbers of defectives.

***Index Terms*—Fingerprinting, traitor tracing, group testing, channel capacities, search problems, compressive sensing.**

I. Introduction

A. Fingerprinting

To protect copyrighted content against unauthorized redistribution, distributors commonly embed watermarks or fingerprints in the content, uniquely linking copies to individual users. If the distributor finds an illegal copy of the content online, he can then extract the watermark from this copy and compare it to the database of watermarks, to determine which user was responsible.

To combat this solution, a group of c pirates may try to form a coalition and perform a collusion attack. By comparing their unique versions of the content, they will detect differences in their copies which must be part of the watermark. They can then try to create a mixed pirate copy, where the resulting watermark matches the watermark of different pirates in different segments of the content, making it hard for the distributor to find the responsible users. The goal of the distributor of the content is to assign the watermarks to the users in such a way that, even if many pirates collude, the pirate copy can still be traced back to the responsible users.

B. Group testing

A different area of research that has received considerable attention in the last few decades is group testing, introduced by Dorfman [24] in the 1940s. Suppose a large population contains a small number c of infected (or defective) items. To identify these items, it is possible to perform group tests: testing a subset of the population will lead to a positive test result if this subset contains at least one defective item, and a negative result otherwise. Since the time to run a single test may be very long, the subsets to test need to be chosen in advance, after which all group tests are performed simultaneously. Then, when the test results come back, the subset of defective items needs to be identified. The goal of the game is to identify these defectives using as few group tests as possible, and with a probability of error as small as possible.

C. Model

The above problems of fingerprinting and group testing can be jointly modeled by the following two-person game between (in terms of fingerprinting) the distributor $\mathcal{D}$ and the adversary $\mathcal{C}$ (the set of colluders, or the set of defectives). Throughout the paper we will mostly use terminology from fingerprinting (i.e. users instead of items, colluders instead of defective items), unless we are specifically dealing with group testing results.

First, there is a universe $\mathcal{U}$ of n users, and the adversary is assigned a random subset of users $\mathcal{C} \subseteq \mathcal{U}$ of size $|\mathcal{C}| = c$. This subset $\mathcal{C}$ is unknown to the distributor (but we assume that the distributor does know the size c of $\mathcal{C}$), and the aim of the game for the distributor is ultimately to discover $\mathcal{C}$. The two-person game consists of three phases: (1) the distributor uses an *encoder* to generate a fingerprinting code, used for assigning versions to users; (2) the colluders employ a *collusion channel* to generate the pirate output from their given code words; and (3) the distributor uses a *decoder* to map the pirate output to a set $\mathcal{C}' \subseteq \mathcal{U}$.

1) Encoder: First, the distributor generates a fingerprinting code $\mathcal{X}$ of n binary code words of length ℓ.[1] The parameter ℓ is referred to as the code length, and the distributor would like ℓ to be as small as possible. For the eventual embedded watermark, we assume that for each segment of the content there are two differently watermarked versions, so the watermark of user j is determined by the ℓ entries in the jth code word of $\mathcal{X}$.

A common restriction on the encoding process is to assume that $\mathcal{X}$ is created by first generating a bias vector $\boldsymbol{P} \in (0,1)^\ell$ (by choosing each entry P_i, for $i = 1, \ldots, \ell$, independently

T. Laarhoven is with the Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands.
E-mail: mail@thijs.com.

Part of the material in this paper will be presented at the 2nd ACM Workshop on Information Hiding and Multimedia Security (Salzburg, Austria, June 2014).

[1] In fingerprinting a common generalization is to assume that the entries of the code words come from an alphabet of size $q \geq 2$, but in this paper we restrict our attention to the binary case $q = 2$.

from a certain distribution f_P), and then generating code words $\boldsymbol{X}_j \in \mathcal{X}$ according to $\mathbb{P}(X_{j,i} = 1) = P_i$. This guarantees that watermarks of different users j are independent, and that watermarks in different positions i are independent. Fingerprinting schemes that satisfy this assumption are sometimes called bias-based schemes, and the encoders in this paper (both for group testing and fingerprinting) are also assumed to belong to this category.

2) Collusion channel: After generating $\mathcal{X}$, the code words are used to select and embed watermarks in the content, and the content is sent out to all users. The colluders then get together, compare their copies, and use a certain collusion channel or pirate attack $\boldsymbol{\Theta}$ to determine the pirate output $\boldsymbol{Y} \in \{0,1\}^\ell$. If the pirate attack behaves symmetrically both in the colluders and in the positions i, then the collusion channel can be modeled by a vector $\boldsymbol{\theta} \in [0,1]^{c+1}$, consisting of entries $\theta_z = \mathbb{P}(Y_i = 1|z)$ (for $z = 0, \dots, c$) indicating the probability of outputting a 1 when the pirates received z ones and $c-z$ zeroes. A further restriction on $\boldsymbol{\theta}$ in fingerprinting is the marking assumption introduced by Boneh and Shaw [11], which says that $\theta_0 = 0$ and $\theta_c = 1$, i.e., if the pirates receive only zeros or ones they have to output this symbol.

3) Decoder: Finally, after the pirate output has been generated and distributed, we assume the distributor intercepts it and applies a decoding algorithm to $\boldsymbol{Y}$, $\mathcal{X}$ and $\boldsymbol{P}$ to compute a set $\mathcal{C}' \subseteq \mathcal{U}$ of accused users. The distributor wins the game if $\mathcal{C}' = \mathcal{C}$ and loses if this is not the case.[2]

Fingerprinting vs. group testing: While the above model is described in fingerprinting terminology, it also covers many common group testing models. The users then correspond to items, the colluders translate to defectives, the code $\mathcal{X}$ corresponds to the group testing matrix X (where $X_{j,i} = 1$ if item j is included in the ith test), and the pirate output corresponds to positive/negative test results. The collusion channel is exactly what separates group testing from fingerprinting: while in fingerprinting it is commonly assumed that this channel is not known or only weakly known to the distributor, in group testing this channel is usually assumed known in advance. This means that there is no malicious adversary in group testing, but only a randomization procedure that determines $\boldsymbol{Y}$. Note also that in (noisy) group testing, the Boneh-Shaw marking assumption may not always hold.

D. Related work

Work on the fingerprinting game described above started in the late 90s, and lower bounds on the code length were established of the order $\ell \propto c \ln n$ [11], until in 2003 Tardos [52] proved a lower bound of the order $\ell \propto c^2 \ln n$ and described a scheme with $\ell = O(c^2 \ln n)$, showing this bound is tight. The lower bound however did not explicitly mention leading constants, so later work on fingerprinting focused on finding this constant. Amiri and Tardos [5] and Huang and Moulin [29]–[32], [44] independently worked on this problem and found that the exact asymptotic lower bound on ℓ is $\ell \sim 2c^2 \ln n$ for large n and c. Huang and Moulin further derived the pirate strategy and encoder achieving this lower bound, which were later used [39], [42], [45] to construct efficient schemes matching these bounds. Most of this work on lower bounds considers the setting with worst-case pirate attacks, and little is known about lower bounds for specific, suboptimal pirate attacks considered in e.g. [9], [15], [28], [38], [42], [45]. It is well known that for suboptimal pirate attacks the required code length may be significantly smaller than the $\ell \sim 2c^2 \ln n$ for arbitrary attacks, but no tight bounds are known.

Research on the group testing problem started much longer ago, and already in 1985 exact asymptotics on the code length for probabilistic schemes were derived as $\ell \sim c \log_2 n$ [48], whereas deterministic schemes require a code length of $\ell \propto c^2 \ln n$ [25], [26]. Later work focused on slight variations of the classical model such as noisy group testing, where a positive result may not always correspond to the presence of a defective item due to 'noise' in the test output [6], [7], [17], [18], [34], [37], [49], and threshold group testing, where the test result may only be positive if sufficiently many defective items are included in the tested subset [1]–[3], [16], [19], [22], [37], [41]. For noisy and threshold group testing, exact asymptotics on the capacities are yet unknown, and so it is not known whether existing constructions are optimal.

E. Contributions and outline

In this paper we extend the work of Huang and Moulin [32] to various fingerprinting and group testing models where $\boldsymbol{\theta}$ is known in advance. We derive the simple capacities (Section II) and the joint capacities (Section III) for these pirate attacks/group testing models, and prove that these results are the exact optima. An overview of these results can be found in Table I. Finally, we discuss the results (Section IV) and mention some directions for future work (Section V).

II. Simple capacities

In simple decoding, *"the receiver makes an innocent/guilty decision on each user independently of the other users, and there lies the simplicity but also the suboptimality of this decoder."* [44, Section 4.3] In other words, a simple decoder bases its decision whether or not to accuse user j only on the jth code word of $\mathcal{X}$, and not on other code words in $\mathcal{X}$. This means that the decoding step will generally be fast but less accurate than when all information available to the decoder (the entire code $\mathcal{X}$) is taken into account.

Huang and Moulin [29]–[32] previously studied simple capacities in the context of fingerprinting, and showed that given a set of allowed collusion channels $\mathcal{P}^c$ (depending only on $\boldsymbol{\theta}$) and a set of allowed encoders $\mathcal{P}^e$, any fingerprinting rate below

$$C^s(\mathcal{P}^e, \mathcal{P}^c) = \max_{f_P \in \mathcal{P}^e} \min_{\boldsymbol{\theta} \in \mathcal{P}^c} \mathbb{E}_P I(p, \boldsymbol{\theta}) \tag{1}$$

is achievable[3], where

$$I(p, \boldsymbol{\theta}) = I(X_1; Y|P = p) \tag{2}$$

[2] In this paper we thus consider the catch-all scenario, where not at least one colluder (the catch-one scenario) but *all* colluders should be found for the distributor to win the game.

[3] Certain conditions on $\mathcal{P}^c$ and $\mathcal{P}^e$ need to be satisfied for this to hold, but we refer the reader to [32] for details.

TABLE I

AN OVERVIEW OF THE CAPACITY RESULTS DERIVED IN THIS PAPER. THE TOP ROWS DESCRIBE PIRATE ATTACKS IN FINGERPRINTING AND THE BOTTOM ROWS CONSIDER VARIOUS DIFFERENT MODELS IN GROUP TESTING. THE RESULTS FOR THE INTERLEAVING ATTACK AND UNKNOWN ATTACKS [32] AND THE JOINT CAPACITY FOR THE CLASSICAL GROUP TESTING MODEL [48] WERE DERIVED BEFORE, WHILE THE OTHER RESULTS ARE NEW.

	Attacks / Models	Simple capacities	(Section II)	Joint capacities	(Section III)
Fingerprinting	$\boldsymbol{\theta}_{\text{int}}$: interleaving attack	$\left(\frac{1}{2\ln 2}\right)/c^2 \approx 0.72/c^2$	(Section II-A1)	$\left(\frac{1}{2\ln 2}\right)/c^2 \approx 0.72/c^2$	(Section III-A1)
	$\boldsymbol{\theta}_{\text{all1}}$: all-1 attack	$(\ln 2)/c \approx 0.69/c$	(Section II-A2)	$(1)/c \approx 1.00/c$	(Section III-A2)
	$\boldsymbol{\theta}_{\text{maj}}$: majority voting	$\left(\frac{1}{\pi\ln 2}\right)/c \approx 0.46/c$	(Section II-A3)	$(1)/c \approx 1.00/c$	(Section III-A3)
	$\boldsymbol{\theta}_{\text{min}}$: minority voting	$(\ln 2)/c \approx 0.69/c$	(Section II-A4)	$(1)/c \approx 1.00/c$	(Section III-A4)
	$\boldsymbol{\theta}_{\text{coin}}$: coin-flip attack	$\left(\frac{1}{4}\ln 2\right)/c \approx 0.17/c$	(Section II-A5)	$\left(\log_2(\frac{5}{4})\right)/c \approx 0.32/c$	(Section III-A5)
	$\mathcal{P}_{\text{mark}}$: unknown attacks	$\left(\frac{1}{2\ln 2}\right)/c^2 \approx 0.72/c^2$	(Section II-A6)	$\left(\frac{1}{2\ln 2}\right)/c^2 \approx 0.72/c^2$	(Section III-A6)
Group testing	$\boldsymbol{\theta}_{\text{all1}}$: classical model	$(\ln 2)/c \approx 0.69/c$	(Section II-B1)	$(1)/c \approx 1.00/c$	(Section III-B1)
	$\boldsymbol{\theta}_{\text{add}}$: additive noise	$(\ln 2 - r)/c \approx 0.69/c$	(Section II-B2)	$(1-\frac{1}{2}h(r))/c \approx 1.00/c$	(Section III-B2)
	$\boldsymbol{\theta}_{\text{dil}}$: dilution noise	$(\ln 2 - O(r\ln r))/c \approx 0.69/c$	(Section II-B3)	$(1-\frac{1}{2}h(r)\ln 2)/c \approx 1.00/c$	(Section III-B3)
	$\boldsymbol{\theta}_{\text{thr}}^{(u)}$: threshold (no gap)	between $0.46/c$ and $0.69/c$	(Section II-B4)	$(1)/c \approx 1.00/c$	(Section III-B4)
	$\boldsymbol{\theta}_{\text{int}}^{(l,u)}$: threshold (int. gap)	between $0.72/c^2$ and $0.69/c$	(Section II-B5)	between $0.72/c^2$ and $1.00/c$	(Section III-B5)
	$\boldsymbol{\theta}_{\text{coin}}^{(l,u)}$: threshold (coin. gap)	between $0.17/c$ and $0.69/c$	(Section II-B5)	between $0.32/c$ and $1.00/c$	(Section III-B5)

is the mutual information between a colluder's symbol X_1 and the pirate output Y in one segment i. In this work we will always let $\mathcal{P}^e$ be the set of all probability distribution functions on $(0,1)$, and we will commonly omit the argument $\mathcal{P}^e$ from C^s. For fixed collusion channels $\mathcal{P}^c = \{\boldsymbol{\theta}\}$, fixing P is optimal [32, Section IV.B], in which case the expression from (1) reduces to

$$C^s(\{\boldsymbol{\theta}\}) = \max_{f_P} \mathbb{E}_P I(p,\theta) = \max_{p\in(0,1)} I(X_1;Y|P=p). \quad (3)$$

With slight abuse of notation we will abbreviate the left hand side as $C^s(\boldsymbol{\theta})$. We will also consider some cases where (part of) the collusion channel is unknown, in which case the capacity includes a minimization over $\boldsymbol{\theta}$ as well:

$$C^s(\mathcal{P}^c) = \max_{f_P} \min_{\boldsymbol{\theta}\in\mathcal{P}^c} \mathbb{E}_P I(X_1;Y|P=p). \quad (4)$$

To study the mutual information payoff function $I(p,\boldsymbol{\theta})$ we will use the following identity [32, Equation (61)]:

$$I(p,\boldsymbol{\theta}) = pd(a_1\|a) + (1-p)d(a_0\|a), \quad (5)$$

where a, a_0, a_1 are defined as

$$a = \sum_{z=0}^{c} \binom{c}{z} p^z (1-p)^{c-z}\theta_z, \quad (6)$$

$$a_0 = \sum_{z=0}^{c-1} \binom{c-1}{z} p^z (1-p)^{c-z-1}\theta_z, \quad (7)$$

$$a_1 = \sum_{z=1}^{c} \binom{c-1}{z-1} p^{z-1} (1-p)^{c-z}\theta_z. \quad (8)$$

Here, $d(\cdot\|\cdot)$ denotes the relative entropy or Kullback-Leibler divergence, defined by $d(\alpha\|\beta) = \alpha\log_2(\frac{\alpha}{\beta}) + (1-\alpha)\log_2(\frac{1-\alpha}{1-\beta})$. Given p and $\boldsymbol{\theta}$, the above formulas allow us to compute the mutual information $I(p,\boldsymbol{\theta})$ explicitly.

For obtaining the simple capacities for various models, we will extensively work with the Kullback-Leibler divergence. In general analyzing this function is not so pretty, but we can simplify some computations with the following Taylor expansion around $\alpha = \beta$:

$$d(\alpha\|\beta) = \frac{(\alpha-\beta)^2}{2\beta(1-\beta)\ln 2}\left(1 + O\left(\frac{|\alpha-\beta|}{\beta(1-\beta)}\right)\right). \quad (9)$$

Intuitively, this says that the divergence is bigger if α and β are further apart, but for α and β both close to 0 or 1 the divergence may blow up as well due to the β and $1-\beta$ in the denominator. In that case we have to be careful and see whether $|\alpha-\beta|$ approaches 0 faster than β or $1-\beta$. A special case of (9) for $\beta = \frac{1}{2}$ and $\alpha \approx \frac{1}{2}$ is

$$d\left(\frac{1}{2}\pm\gamma\middle\|\frac{1}{2}\right) = \frac{2\gamma^2}{\ln 2} + O(\gamma^4). \quad (10)$$

Finally, if $\alpha = \frac{1}{2}$ and $\beta \approx \frac{1}{2}$, we can rewrite $d(\alpha\|\beta)$ as

$$d\left(\frac{1}{2}\middle\|\frac{1}{2}(1\pm\gamma)\right) = \frac{1}{2}d\left(1\|1-\gamma^2\right), \quad (11)$$

and regardless of α and β, we always have $d(\alpha\|\beta) = d(1-\alpha\|1-\beta)$.

A. Fingerprinting

We will study the simple capacities for five commonly considered fingerprinting attacks, and for completeness also mention the result of Huang and Moulin regarding the simple capacity for unknown attacks.

1) Interleaving attack: The interleaving attack in fingerprinting (considered in e.g. [9], [15], [28], [32], [38], [45]) is characterized by the coalition choosing one of its members at random, and outputting his symbol. Given z members with a 1 and $c-z$ members with a 0, the probability of outputting a 1 is then equal to $\frac{z}{c}$, regardless of z and c:

$$(\boldsymbol{\theta}_{\text{int}})_z = \frac{z}{c}. \qquad (0 \le z \le c) \quad (12)$$

This attack is known to be one of the strongest pirate attacks, and the capacity is proportional to $\frac{1}{c^2}$. The exact asymptotics of

the simple capacity for the interleaving attack were previously derived by Huang and Moulin.

Proposition 1: [32, Theorem 6] The simple capacity for the interleaving attack is:

$$C^s(\boldsymbol{\theta}_{\text{int}}) = \frac{1}{2c^2 \ln 2} + O\left(\frac{1}{c^4}\right) \approx \frac{0.72}{c^2}, \tag{13}$$

and the maximizing value of p is $p^s_{\text{int}} = \frac{1}{2}$.

2) All-1 attack: Another commonly considered attack is the all-1 attack, where pirates output a 1 whenever they can [15], [38], [42], [45]. Due to the marking assumption they are forced to output a 0 when they did not receive any ones, but otherwise a coalition using the all-1 attack will always output a 1:

$$(\boldsymbol{\theta}_{\text{all1}})_z = \begin{cases} 0 & \text{if } z = 0; \\ 1 & \text{if } z > 0. \end{cases} \tag{14}$$

We will show below that this attack is significantly weaker than the interleaving attack.

Proposition 2: The simple capacity and the corresponding maximizing value of p for the all-1 attack are:

$$C^s(\boldsymbol{\theta}_{\text{all1}}) = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.69}{c}, \tag{15}$$

$$p^s_{\text{all1}} = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.69}{c}. \tag{16}$$

Proof: First, consider a, a_0 and a_1. Using $\theta_z = 0$ if $z = 0$ and $\theta_z = 1$ otherwise, we get

$$a = \sum_{z=0}^{c} \binom{c}{z} p^z (1-p)^{c-z} \theta_z = 1 - (1-p)^c. \tag{17}$$

Working out a_0 and a_1 in a similar way, we get $a_0 = 1 - (1-p)^{c-1}$ and $a_1 = 1$. For ease of notation, let us write $s = (1-p)^c$ and $I(p) = I(p, \boldsymbol{\theta}_{\text{all1}})$, so that we get

$$I(p) = pd\left(1 \| 1-s\right) + (1-p)d\left(\frac{s}{1-p}\middle\| s\right). \tag{18}$$

Now, consider the second term. For large c, we argue that this term is small, i.e. of the order $O(\frac{1}{c^2})$, regardless of p:

$$\begin{aligned}
&(1-p)d\left(\frac{s}{1-p} \;\middle\|\; s\right) \\
&= -s\log_2(1-p) + (1-p-s)\log_2\left(1 - \frac{ps}{(1-p)(1-s)}\right) \\
&\overset{(a)}{=} -s\log_2(1-p) + \frac{1-p-s}{\ln 2}\left[\frac{-ps}{(1-p)(1-s)} + O\left(\frac{1}{c^2}\right)\right] \\
&\overset{(b)}{=} +\frac{ps}{\ln 2} - \frac{ps}{\ln 2}\left[1 - \frac{ps}{(1-p)(1-s)} + O\left(\frac{1}{c^2}\right)\right] \\
&\overset{(c)}{=} +\frac{ps}{\ln 2} - \frac{ps}{\ln 2} + O\left(\frac{1}{c^2}\right) = O\left(\frac{1}{c^2}\right).
\end{aligned}$$

Here (a) follows from $\frac{ps}{(1-p)(1-s)} = O(\frac{1}{c})$ and $ps = O(\frac{1}{c})$ for all p, (b) follows from $p^2 s = O(\frac{1}{c^2})$, and (c) follows from $\frac{p^2 s^2}{(1-p)(1-s)} = O(\frac{1}{c^2})$ and $p^2 s = O(\frac{1}{c^2})$ for arbitrary p. So we are now left with:

$$I(p) = -p\log_2(1-s) + O\left(\frac{1}{c^2}\right). \tag{19}$$

For p to be a global maximum we need either that $I'(p) = 0$ or p should be one of the end-points 0 or 1. For $p \to 0, 1$ we get $I(p) \to 0$, so we need to find a value $p \in (0,1)$ with $I'(p) = 0$. Writing out the remaining term and differentiating, this condition is equivalent to

$$\frac{cps}{(1-p)(1-s)} = -\ln(1-s). \tag{20}$$

Since the left hand side is $O(1)$ regardless of p, the right hand side must be too, so $s = 1 - o(1)$ is excluded. To exclude the case $s = o(1)$ we rewrite (20) to get

$$\frac{cp}{1-p} = \frac{1-s}{s}\ln\left(\frac{1}{1-s}\right). \tag{21}$$

Now if $s = o(1)$ then the right hand side becomes $1 - o(1)$, which implies in the left hand side that $p = \frac{1}{c} - o(\frac{1}{c})$, which implies that $s \neq o(1)$, contradicting our assumption that $s = o(1)$. So for large c a maximum can only occur at $o(1) < s < 1 - o(1)$. Suppose that $s(c) \to s^* \in (0,1)$ for $c \to \infty$, with $s^* \neq s^*(c)$ not depending on c. Then $p(c) \to p^* = \frac{-1}{c}\ln s^*$, so the condition on p and s is then asymptotically equivalent to:

$$s^* \ln s^* = (1-s^*)\ln(1-s^*) + O\left(\frac{1}{c}\right). \tag{22}$$

This has a unique solution at $s^* = \frac{1}{2} + O(\frac{1}{c})$, leading to the given values of p^s_{all1} and $C^s(\boldsymbol{\theta}_{\text{all1}})$. ∎

In terms of code lengths, this means that any simple decoding algorithm for the all-1 attack requires an asymptotic number of fingerprint positions of at least $\ell \sim \frac{1}{\ln 2} c \log_2 n \approx 2.08 c \ln n$ for large n. This seems to contradict earlier results of [37], [38], which suggested that under a certain Gaussian assumption, only $\ell \sim 2c \ln n$ tests are required. This apparent contradiction is caused by the fact that the Gaussian assumption in that paper is not correct in the regime of small p, for which those results were derived. Rigorous analysis of the scores in [37], [38] shows that with that scheme, an asymptotic code length of about $\ell \approx 3c \ln n$ is sufficient when $p \sim \frac{1}{c}\ln(2)$, which is well above the lower bound obtained above.

3) Majority voting: The majority voting attack [9], [15], [28], [38], [43], [45] is characterized by the pirates choosing the symbol they have seen the most often. To avoid ambiguity, we will assume c is odd, in which case the attack is given by

$$(\boldsymbol{\theta}_{\text{maj}})_z = \begin{cases} 0 & \text{if } z < \frac{c}{2}; \\ 1 & \text{if } z > \frac{c}{2}. \end{cases} \tag{23}$$

For this attack we obtain the following result.

Proposition 3: For the majority voting attack, the simple capacity is

$$C^s(\boldsymbol{\theta}_{\text{maj}}) = \frac{1}{\pi c \ln 2} + O\left(\frac{1}{c^2}\right) \approx \frac{0.46}{c}, \tag{24}$$

and the maximizing value of p is $p^s_{\text{maj}} = \frac{1}{2}$.

Proof: As mentioned before, to avoid ambiguity we focus on the case where $c = 2c' + 1$ is odd, and due to symmetry

w.l.o.g. we may assume that $p \le \frac{1}{2}$. First, we have:

$$a = \sum_{z=c'+1}^{2c'+1} \binom{2c'+1}{z} p^z (1-p)^{2c'+1-z}, \tag{25}$$

and a_0 and a_1 satisfy $a_0 = a + pu$ and $a_1 = a - (1-p)u$, where $u = \binom{2c'}{c'} p^{c'}(1-p)^{c'}$. Now if $p = O(\frac{1}{c})$, then a_1 and a_0 quickly approach 0 leading to $I(p) = o(\frac{1}{c})$. For the remaining case $p = \omega(\frac{1}{c})$, expanding a using Sanov's theorem [21, Theorem 11.4.1] we get

$$a \sim \exp\left[(2c'+1)\ln(2) d\left(\tfrac{1}{2}\|p\right)\right] \tag{26}$$

$$\sim p^{c'+\frac{1}{2}}(1-p)^{c'+\frac{1}{2}} 2^{2c'+1}. \tag{27}$$

Using Stirling's formula for the central binomial coefficient in u, we obtain

$$u = \binom{2c'}{c'} p^{c'}(1-p)^{c'} \sim \frac{2^{2c'} p^{c'}(1-p)^{c'}}{\sqrt{\pi c'}}. \tag{28}$$

As a consequence, $\frac{u}{a} = o(1)$, and using (9) we get

$$d(a_0\|a) \sim \frac{p^2 u^2}{2\ln 2 a(1-a)}, \tag{29}$$

$$d(a_1\|a) \sim \frac{(1-p)^2 u^2}{2\ln 2 a(1-a)}. \tag{30}$$

Combining these expressions, we get

$$I(p) = p d(a_1\|a) + (1-p) d(a_0\|a) \tag{31}$$

$$\sim \frac{2^{4c'} p^{2c'+1}(1-p)^{2c'+1}}{2\pi c' a(1-a)\ln 2}. \tag{32}$$

To see that this has a maximum at $p = \frac{1}{2}$, writing out the inverse of the above expression (ignoring constants) we see that, in terms of p,

$$\frac{1}{I(p)} \propto \sum_{z_1,z_2=0}^{c'} \binom{2c'+1}{z_1}\binom{2c'+1}{z_2}\left(\frac{p}{1-p}\right)^{z_1-z_2} \tag{33}$$

$$= C_1 + \sum_{z_1<z_2} C_2 \left[\left(\frac{p}{1-p}\right)^{z_2-z_1} + \left(\frac{1-p}{p}\right)^{z_2-z_1}\right] \tag{34}$$

$$= C_1 + \sum_{z_1<z_2} C_2 \left[2\cosh((z_2-z_1)\ln x)\right], \tag{35}$$

where $x = \frac{1-p}{p} > 1$ for $p < \frac{1}{2}$ and $x = 1$ if $p = \frac{1}{2}$, and C_1, C_2 are expressions that do not depend on p. The function between square brackets is positive and increasing in x for $x \ge 1$, so it has a global minimum at $x = 1$, corresponding to $p = \frac{1}{2}$. So the maximum for $I(p)$ is attained at $p = \frac{1}{2}$, in which case u satisfies

$$u = \frac{1}{\sqrt{\pi c/2}}\left(1 + O\left(\frac{1}{c}\right)\right). \tag{36}$$

To get exact asymptotics for $I(\frac{1}{2})$, we return to the expression for $I(p)$ of (9). Since from (25) it follows that $a = \frac{1}{2}$, and both terms are identical, we obtain:

$$I\left(\frac{1}{2}\right) = d\left(\frac{1}{2} + \frac{1}{\sqrt{2\pi c}}\left[1 + O\left(\frac{1}{c}\right)\right] \Big\| \frac{1}{2}\right). \tag{37}$$

Using (10) the result then follows. ∎

This result matches the bounds obtained in [37], [38], which showed that with an almost trivial decoding algorithm one can asymptotically achieve a code length of $\ell \sim \pi c \ln n$ for large n and c. The construction of [37], [38] is thus capacity-achieving.

4) Minority voting: As the name suggests, when pirates use the minority voting attack [9], [15], [28], [38], [45], they output the symbol they have received the least often. Due to the marking assumption they are not able to output symbols they have not received, so in the binary setting the attack is defined as follows. Again, we will assume that c is odd.

$$(\boldsymbol{\theta}_{\text{min}})_z = \begin{cases} 0 & \text{if } z = 0 \text{ or } \frac{c}{2} < z < c; \\ 1 & \text{if } z = c \text{ or } 0 < z < \frac{c}{2}. \end{cases} \tag{38}$$

As shown below, this attack has the same simple capacity as the all-1 attack.

Proposition 4: The simple capacity and the corresponding optimal value of p for the minority voting attack are:

$$C^s(\boldsymbol{\theta}_{\text{min}}) = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.69}{c}, \tag{39}$$

$$p^s_{\text{min}} = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.69}{c}. \tag{40}$$

Proof: In this case the function $I(p)$ is symmetric around $p = \frac{1}{2}$, so w.l.o.g. we may assume $p \le \frac{1}{2}$. For small values of p, minority voting is equivalent to the all-1 attack up to negligible order terms, while for $p \approx \frac{1}{2}$ the attack is very similar to majority voting by $\boldsymbol{\theta}_{\text{min}} \approx 1 - \boldsymbol{\theta}_{\text{maj}}$. This means that for small p the mutual information payoff will be equivalent to that of the all-1 attack, while for $p \approx \frac{1}{2}$ we get the same values as for majority voting. Since the simple capacity for the all-1 attack is higher than for majority voting, the distributor should choose p close to p^s_{all1}, leading to the result. ∎

5) Coin-flip attack: Instead of choosing a pirate at random and outputting his symbol (the interleaving attack), the pirates may also decide to choose a symbol at random from their set of received symbols, without paying attention to how often they received each symbol [9], [28], [38], [45]. In other words, when a coalition receives both symbols, they let a fair coin-flip decide which symbol to output. This means that the collusion channel satisfies:

$$(\boldsymbol{\theta}_{\text{coin}})_z = \begin{cases} 0 & \text{if } z = 0; \\ \frac{1}{2} & \text{if } 0 < z < c; \\ 1 & \text{if } z = c. \end{cases} \tag{41}$$

This pirate attack is weaker than the interleaving attack, but stronger than the other pirate attacks considered above.

Proposition 5: For the coin-flip attack, the simple capacity and the corresponding maximizing value of p are:

$$C^s(\boldsymbol{\theta}_{\text{coin}}) = \frac{\ln 2}{4c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.17}{c}, \tag{42}$$

$$p^s_{\text{coin}} = \frac{\ln 2}{2c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.35}{c}. \tag{43}$$

Proof: Since $I(p)$ is symmetric around $p = \frac{1}{2}$, let us assume w.l.o.g. that $p \leq \frac{1}{2}$. For a, a_0 and a_1 we obtain:

$$a = \frac{1}{2}\left(1 + p^c - (1-p)^c\right), \tag{44}$$

$$a_0 = \frac{1}{2}\left(1 - (1-p)^{c-1}\right), \tag{45}$$

$$a_1 = \frac{1}{2}\left(1 + p^{c-1}\right). \tag{46}$$

So for the mutual information, we obtain

$$I(p) = pd\left(\frac{1}{2}\left(1+p^{c-1}\right) \middle\| \frac{1}{2}\left(1+p^c-(1-p)^c\right)\right) \tag{47}$$

$$+ (1-p)d\left(\frac{1}{2}\left(1-(1-p)^{c-1}\right) \middle\| \frac{1}{2}\left(1+p^c-(1-p)^c\right)\right). \tag{48}$$

For $p \leq \frac{1}{2}$, the terms p^c and p^{c-1} are negligible, so up to small order terms, we get

$$I(p) = pd\left(\frac{1}{2}\middle\| \frac{1}{2}\left(1-(1-p)^c\right)\right) \tag{49}$$

$$+ (1-p)d\left(\frac{1}{2}\left(1-(1-p)^{c-1}\right) \middle\| \frac{1}{2}\left(1-(1-p)^c\right)\right). \tag{50}$$

Similar to the proof of the all-1 attack, the second term is $O(\frac{1}{c^2})$, while using (11) we can rewrite the first term to a recognizable form:

$$I(p) = \frac{1}{2}\left[-p\log\left(1-(1-p)^{2c}\right)\right] + O\left(\frac{1}{c^2}\right). \tag{51}$$

The term between square brackets is exactly the dominating term for the simple capacity of the all-1 attack for $c' = 2c$. In other words:

$$I_c(p, \boldsymbol{\theta}_{\text{coin}}) = \frac{1}{2} I_{2c}(p, \boldsymbol{\theta}_{\text{all1}}) + O\left(\frac{1}{c^2}\right). \tag{52}$$

Using Proposition 2, the result follows. ∎

For this attack, the result in [38] was also too optimistic due to the incorrect Gaussian assumption. Any simple decoder must have a code length of at least $\ell \sim \frac{4}{\ln 2} c \log_2 n \approx 8.33 c \ln n$, while the result in [38] suggests that a code length of $\ell \sim 4c \ln n$ suffices under a certain Gaussian assumption. Again, the Gaussian assumption is to blame, and since the optimal value of p is even smaller here than for the all-1 attack, the error of [38] is even bigger here.

6) Unknown attacks: Finally, the most often studied setting in fingerprinting is the scenario where the pirate attack is not known to the distributor. Due to the marking assumption the distributor does know that $\theta_0 = 0$ and $\theta_c = 1$, but otherwise no assumptions are made on the pirate strategy. The set of allowed attacks can then be described as

$$\mathcal{P}_{\text{mark}} = \{\boldsymbol{\theta} \in [0,1]^{c+1} \mid \theta_0 = 0, \theta_c = 1\}. \tag{53}$$

Huang and Moulin solved the related max-min game for large c, and found the asymptotic optimal encoder and collusion channel leading to the saddle point solution.

Proposition 6: [32, Theorem 6, Corollary 7] The simple capacity for the uninformed fingerprinting game is

$$C^s(\mathcal{P}_{\text{mark}}) = \frac{1}{2c^2 \ln 2} + O\left(\frac{1}{c^3}\right) \approx \frac{0.72}{c^2}, \tag{54}$$

and the optimizing encoder f_P and collusion channel $\boldsymbol{\theta}$ achieving this bound for large c are the arcsine distribution, defined by

$$f_P^*(p) = \frac{1}{\pi\sqrt{p(1-p)}}, \qquad (p \in (0,1)) \tag{55}$$

and the interleaving attack $\boldsymbol{\theta}_{\text{int}}$.

B. Group testing

For group testing, we will study five different models: the classical (noiseless) model, the models with additive noise and dilution noise, and threshold group testing with and without gaps. Other models where the test result Y depends only on the tally Z may be analyzed in a similar fashion.

1) Classical model: In the classical model, the outcome of a group test is positive iff at least one defective was present in the tested pool. This model is equivalent to the all-1 attack in fingerprinting, as was previously noted in e.g. [37], [42], [51]. This immediately leads to the following result.

Corollary 1: For the classical group testing model, the simple informed capacity and the corresponding optimal value of p are:

$$C^s(\boldsymbol{\theta}_{\text{all1}}) = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.69}{c}, \tag{56}$$

$$p^s_{\text{all1}} = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.69}{c}. \tag{57}$$

In terms of group testing algorithms, this means that any simple decoding algorithm for c defectives and n total items requires an asymptotic number of group tests ℓ of at least

$$\ell \sim \frac{c \log_2 n}{\ln 2} \approx 1.44\, c \log_2 n \approx 2.08\, c \ln n, \tag{58}$$

where the asymptotics are for $n \to \infty$. This improves upon the known lower bound for joint decoders of $\ell \geq c \log_2 n$ for large n [48], and this shows that the algorithm of Chan et al. [12] (which achieves a code length of $\ell \sim e \ln n$) is suboptimal. The related paper [40] shows how this bound can actually be achieved with efficient simple decoders.

2) Additive noise: The classical group testing model is sometimes considered to be too optimistic, as the outcome of the group tests may not always be accurate. One 'noisy' variant of the classical model that is sometimes considered in the literature is the additive noise model [7], [12], [18], [49], where a test result may even be positive (with some small probability r) if there were no defectives in the tested group. This corresponds to the following channel $\boldsymbol{\theta}_{\text{add}}$:

$$(\boldsymbol{\theta}_{\text{add}})_z = \begin{cases} r & \text{if } z = 0; \\ 1 & \text{if } z > 0. \end{cases} \tag{59}$$

For small r we do not expect the simple capacity or the optimal choice of p to change drastically compared to the classical model, and the following analysis confirms this.

Proposition 7: For the additive noise model with parameter r, the simple capacity and the maximizing value of p are:

$$C^s(\boldsymbol{\theta}_{\text{add}}) = \frac{\ln 2}{c}\left(1 - \frac{r}{\ln 2} + O(r^2)\right) + O\left(\frac{1}{c^2}\right), \tag{60}$$

$$p^s_{\text{add}} = \frac{\ln 2}{c}\left(1 + \frac{r(2\ln 2 - 1)}{2\ln 2(1-\ln 2)} + O(r^2)\right) + O\left(\frac{1}{c^2}\right). \tag{61}$$

Proof: Working out a, a_0 and a_1, and substituting them into $I(p) = pd(a_1\|a) + (1-p)d(a_0\|a)$, we obtain

$$I(p) = pd(1\|1-(1-p)^c(1-r)) \tag{62}$$

$$+ (1-p)d((1-p)^{c-1}(1-r)\|(1-p)^c(1-r)). \tag{63}$$

For similar reasons as for the all-1 attack, for small values of r the second term is $O(\frac{1}{c^2})$ while the first term is $\Theta(\frac{1}{c})$ and dominates the expression for large c. This means that for small r we have

$$I(p) = -p\log_2(1-(1-p)^c(1-r)) + O\left(\frac{1}{c^2}\right). \tag{64}$$

To find the maximum we take the derivative with respect to p and set it equal to 0 to obtain

$$\ln(1-(1-p)^c(1-r)) = -\frac{cp}{1-p}\cdot\frac{(1-p)^c(1-r)}{1-(1-p)^c(1-r)}. \tag{65}$$

For small r, the above expression is very close to the one we had for the all-1 attack, and again the optimal value of p is close to $\frac{\ln 2}{c}$. Writing $s = (1-p)^c(1-r)$, so that $p = \frac{-1}{c}\ln(\frac{s}{1-r}) + O(\frac{1}{c^2})$ and $1-p = 1-O(\frac{1}{c})$, the above expression reduces to

$$\ln(1-s) = \ln\left(\frac{s}{1-r}\right)\cdot\frac{s}{1-s} + O\left(\frac{1}{c}\right). \tag{66}$$

For small r, this means that $s \approx \frac{1}{2}$, so suppose $s = \frac{1}{2}(1+\varepsilon)$. Filling this in in the above equation, Tayloring around $\varepsilon = 0$, and disregarding terms of the order $\varepsilon^2, r^2, \varepsilon r$, we get

$$-\ln 2 - \varepsilon = (-\ln 2 + r + \varepsilon)(1+2\varepsilon). \tag{67}$$

Rearranging the terms, this leads to

$$\varepsilon = -\frac{r}{2(1-\ln 2)} + O(r^2). \tag{68}$$

Substituting ε into s and solving for p, we get

$$p = -\frac{1}{c}\ln\left(\frac{1}{2}\cdot\frac{1-\frac{r}{2(1-\ln 2)}}{1-r}\right) + O\left(\frac{1}{c^2}\right) \tag{69}$$

$$= \frac{\ln 2}{c} + \frac{r}{c}\cdot\frac{2\ln 2 - 1}{2-2\ln 2} + O\left(\frac{r^2}{c} + \frac{1}{c^2}\right), \tag{70}$$

and for the capacity we get

$$I(p) = -\frac{p}{\ln 2}\ln(1-s) \tag{71}$$

$$= \left[-\frac{1}{c} + \frac{r}{c\ln 2}\cdot\frac{2\ln 2 - 1}{2-2\ln 2}\right]\left[-\ln 2 + \frac{r}{c}\cdot\frac{1}{2-2\ln 2}\right] \tag{72}$$

$$= \frac{\ln 2}{c}\left(1 - \frac{r}{\ln 2} + O(r^2)\right) + O\left(\frac{1}{c^2}\right). \tag{73}$$

■

For small values of r, one should therefore take p to be slightly smaller than $\frac{1}{c}\ln 2$, and the capacity will be slightly lower than in the classical model.

3) Dilution noise: Another commonly considered noisy group testing model is the dilution noise model [7], [17], [18], [34], [49], where the probability of a positive test outcome depends on the number of defectives in the tested pool. More precisely, $\boldsymbol{\theta}_{\text{dil}}$ is defined as follows:

$$(\boldsymbol{\theta}_{\text{dil}})_z = \begin{cases} 0 & \text{if } z = 0; \\ 1 - r^z & \text{if } z > 0. \end{cases} \tag{74}$$

Again, for small r this model is close to the traditional group testing model, so both the capacity and the optimal value of p are close to the values of Proposition 2.

Proposition 8: For the dilution noise model with parameter r, neglecting terms of the order c^{-2} and r^2, the simple capacity and the corresponding optimal value of p are:

$$C^s(\boldsymbol{\theta}_{\text{dil}}) = \frac{\ln 2}{c}\left(1 + \frac{r\ln r}{2\ln 2} - \frac{r(1-\ln 2)}{2\ln 2} + O(r^2\ln r)\right) + O\left(\frac{1}{c^2}\right) \tag{75}$$

$$p^s_{\text{dil}} = \frac{\ln 2}{c}\left(1 + \frac{r\ln r}{4\ln 2} + \frac{r(-3(\ln 2)^2 + 5\ln 2 - 1)}{4\ln 2(1-\ln 2)}\right. \tag{76}$$

$$\left. + O(r^2\ln r)\right) + O\left(\frac{1}{c^2}\right). \tag{77}$$

Proof: For a, a_0 and a_1 we get

$$a = 1 - (1-p+pr)^c, \tag{78}$$

$$a_0 = 1 - (1-p+pr)^{c-1}, \tag{79}$$

$$a_1 = 1 - r(1-p+pr)^{c-1}, \tag{80}$$

so letting $s = (1-p+pr)^c$, the mutual information satisfies

$$I(p) = pd\left(\frac{rs}{1-p+pr}\|s\right) + (1-p)d\left(\frac{s}{1-p+pr}\|s\right). \tag{81}$$

For small r, the second term is again small. So expanding the left term, knowing that $p = \Theta(\frac{1}{c})$, we obtain:

$$I(p) = \frac{p}{\ln 2}\left(rs\ln r + (1-rs)\ln\left(\frac{1-rs}{1-s}\right)\right). \tag{82}$$

Writing $p = \frac{\ln 2}{c}(1+\varepsilon)$, we can Taylor s and rs (disregarding terms of the order $r^2, r\varepsilon^2, \varepsilon^3, \frac{1}{c}$) to obtain

$$s = \frac{1}{2}\left(1 - \varepsilon\ln 2 + r\ln 2 + \varepsilon r\ln 2(1-\ln 2) + \frac{\varepsilon^2}{2}(\ln 2)^2\right). \tag{83}$$

This means that up to small order terms, we get $rs = \frac{1}{2}(r - \varepsilon r\ln 2)$. Plugging these into the expression for $I(p)$, we eventually get

$$I(p) = \frac{\ln 2}{c}\left(1 + r\left(\frac{\ln r - 1 + \ln 2}{2\ln 2}\right) + \varepsilon^2(\ln 2 - 1)\right. \tag{84}$$

$$\left. + \varepsilon r\left(\frac{\ln r(1-\ln 2) - 3(\ln 2)^2 + 5\ln 2 - 1}{2\ln 2}\right) + O(\dots)\right). \tag{85}$$

This immediately leads to the given expression for the capacity by disregarding small terms, while differentiating with respect to ε and setting equal to 0 leads to

$$\varepsilon = \left(\frac{\ln r(1-\ln 2) - 1 + 5\ln 2 - 3(\ln 2)^2}{4\ln 2(1-\ln 2)} \right) r + O(r^2). \tag{86}$$

This leads to the given expression for p. ∎

4) Threshold without gaps: Besides accounting for possible mistakes in the test results (noisy group testing), models have also been considered to account for sensitivity in detecting positive items. In threshold group testing [1]–[3], [14], [19], [22], [41], it is assumed that if the number of defectives z in the tested pool is at most l then the test comes back negative, and if z is at least u then the test result is always positive. For the case $u = l + 1$, which we will refer to as threshold group testing without a gap (where $g = u - l - 1$ is the gap size), this completely determines the model:

$$(\boldsymbol{\theta}_{\text{thr}}^{(u)})_z = \begin{cases} 0 & \text{if } z < u; \\ 1 & \text{if } z \geq u. \end{cases} \tag{87}$$

Although simple to state, even for small u and c finding the simple capacity and optimal choice of p analytically seems very hard, if not impossible. We can intuitively see how the capacity will roughly behave though, since we know that:

- The case $u = 1$ corresponds to $\boldsymbol{\theta}_{\text{thr}}^{(u)} = \boldsymbol{\theta}_{\text{all1}}$, for which $p = \frac{\ln 2}{c}$ and $I \approx \frac{\ln 2}{c} \approx \frac{1.44}{c}$ are optimal.
- The case $u = \frac{c+1}{2}$ corresponds to $\boldsymbol{\theta}_{\text{thr}}^{(u)} = \boldsymbol{\theta}_{\text{maj}}$, for which $p = \frac{1}{2}$ and $I = \frac{1}{\pi c \ln 2} \approx \frac{0.46}{c}$ are optimal.

For values of u between 1 and $\frac{c}{2}$, we expect the capacity to decrease as u increases, and the optimal value p is expected to be close to $\frac{u}{c}$.

Numerical evidence supports this intuition, as it shows that the capacity strictly decreases from $u = 1$ up to $u = \frac{c+1}{2}$, and that the optimal values of p are almost evenly spaced for $u = 1$ up to $u = \frac{c}{2}$. The capacity quickly drops at small values of u, i.e., the gap between $C^s(\boldsymbol{\theta}_{\text{thr}}^{(1)})$ and $C^s(\boldsymbol{\theta}_{\text{thr}}^{(2)})$ is bigger than the gap between $C^s(\boldsymbol{\theta}_{\text{thr}}^{(2)})$ and $C^s(\boldsymbol{\theta}_{\text{thr}}^{(13)})$ for $c = 25$.

5) Threshold with gaps: An even harder case to deal with is threshold group testing with $g = u - l - 1 > 0$, which we will refer to as threshold group testing with a gap. If $u > l+1$, then the model is not yet defined properly, as we do not know what θ_z is for $l + 1 \leq z \leq u - 1$. Different models were considered to capture the behavior of the outcome of the test results in these gaps, such as: [14]

- The test outcome is uniformly random:

$$(\boldsymbol{\theta}_{\text{coin}}^{(l,u)})_z = \begin{cases} 0 & \text{if } z \leq l; \\ \frac{1}{2} & \text{if } l < z < u; \\ 1 & \text{if } z \geq u. \end{cases} \tag{88}$$

- The probability of a positive result increases linearly:

$$(\boldsymbol{\theta}_{\text{int}}^{(l,u)})_z = \begin{cases} 0 & \text{if } z \leq l; \\ \frac{z-l}{u-l} & \text{if } l < z < u; \\ 1 & \text{if } z \geq u. \end{cases} \tag{89}$$

- We simply do not know what the test outcome will be.

Note that $\boldsymbol{\theta}_{\text{coin}}^{(0,c)} = \boldsymbol{\theta}_{\text{coin}}$ and $\boldsymbol{\theta}_{\text{int}}^{(0,c)} = \boldsymbol{\theta}_{\text{int}}$, so these models can be seen as generalizations of the corresponding attacks in fingerprinting. Also note that $\boldsymbol{\theta}_{\text{coin}}^{(u-1,u)} = \boldsymbol{\theta}_{\text{int}}^{(u-1,u)} = \boldsymbol{\theta}_{\text{thr}}^{(u)}$.

Regardless of the gap model, for arbitrary l and u these models all seem hard to analyze exactly. Using results obtained previously, we can however try to 'interpolate' the results to get somewhat decent estimates. For instance, for the first model we can interpolate between the results for threshold group testing without a gap (Section II-B4) and the coin-flip attack (Section II-A5) to get upper and lower bounds on the simple capacity. For the second case, we can interpolate between threshold group testing without a gap (Section II-B4) and the interleaving attack (Section II-A1) to get an idea how the capacity and the optimal value of p scale.

To verify this intuition, Figure 1 shows a density plot of the capacities (multiplied by c) for both the coin-flip gap model and the interleaving gap model. These plots are based on numerics for $c = 25$, but already show some trends. For instance, there are sharp peaks in the lower left and upper right corner; even when moving on the diagonal, the capacity quickly drops when leaving the corners. The capacities further take their maxima on and near the diagonal. In the coin-flip gap model, the capacity quickly converges to its minimum at $g = c$ as the gap size increases, while this takes longer for the interleaving gap model. Finally, from Sections II-A1, II-A2, II-A3, and II-A5, we know exactly how the corners and center of each plot behave asymptotically, so we have a decent idea how the capacity scales for large c and arbitrary values of l and u.

III. Joint capacities

Where a simple decoder bases its decision to accuse user j only on the jth code word of $\mathcal{X}$ (and not on other code words), a joint decoder is allowed to use all information available to make a more informed decision. In particular, the whole code $\mathcal{X}$ may be taken into account. Huang and Moulin [29]–[32] previously studied joint capacities as well, and showed that given a set of allowed collusion channels $\mathcal{P}^c$ (depending only on $\boldsymbol{\theta}$) and a set of allowed encoders $\mathcal{P}^e$, any fingerprinting rate below

$$C^j(\mathcal{P}^e, \mathcal{P}^c) = \max_{f_P \in \mathcal{P}^e} \min_{\boldsymbol{\theta} \in \mathcal{P}^c} \mathbb{E}_P I(p, \boldsymbol{\theta}) \tag{90}$$

is achievable, where

$$I(p, \boldsymbol{\theta}) = \frac{1}{c} I(X_1, \ldots, X_c; Y | P = p) \tag{91}$$

is the mutual information between all colluder symbols $X_1, \ldots, X_c$ and the pirate output Y in one segment i. Note that from the assumption that Y only depends on $X_1, \ldots, X_c$ through $\boldsymbol{\theta}$, it follows that $I(X_1, \ldots, X_c; Y|P = p) = I(Z; Y|P = p)$, where $Z = \sum_{i=1}^{c} X_i$. To study the payoff function $I(p, \boldsymbol{\theta}) = I(Z; Y|P = p)$, we will use the following identity [32, Equation (59)]:

$$I(p, \boldsymbol{\theta}) = \frac{1}{c} \left[h(a) - a_h \right] \tag{92}$$

$$\text{with } a_h = \sum_{z=0}^{c} \binom{c}{z} p^z (1-p)^{c-z} h(\theta_z). \tag{93}$$

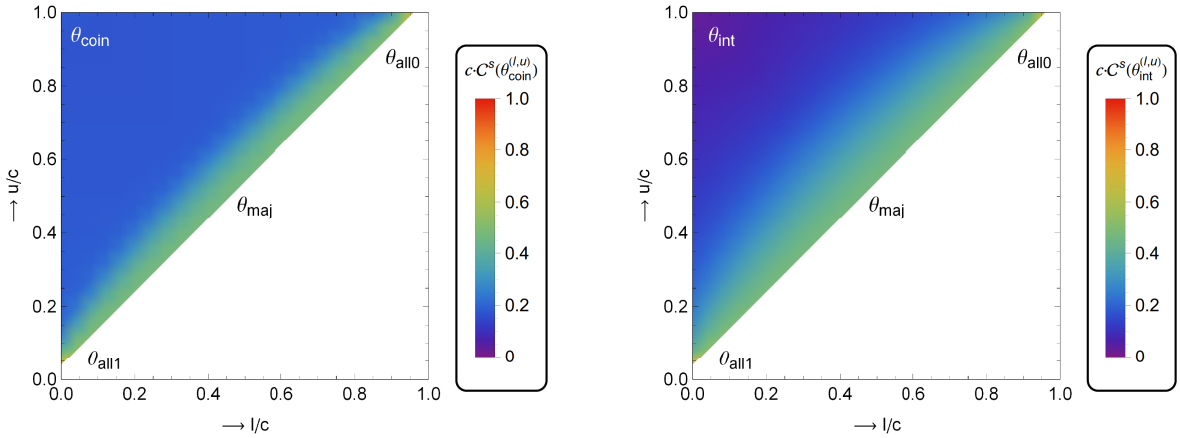

(a) Simple capacity in the coin-flip gap model

(b) Simple capacity in the interleaving gap model

Fig. 1. The simple capacity for threshold group testing for different values of l and u, when there is (a) a coin-flip gap or (b) an interleaving gap. The three corners correspond to the all-1, all-0 and coin-flip or interleaving attack, and the centers of the graphs correspond to the majority voting attack in fingerprinting. The capacity is maximal in the lower left and upper right corner, for which $c \cdot C^j(\boldsymbol{\theta}) \sim \ln 2 \approx 0.69$.

Here $h(\cdot)$ denotes the binary entropy function, defined by $h(\alpha) = -\alpha \log_2 \alpha - (1-\alpha)\log_2(1-\alpha)$. Given p and $\boldsymbol{\theta}$, this allows us to compute $I(p, \boldsymbol{\theta})$ explicitly. In the analysis of specific models $\boldsymbol{\theta}$, we will again commonly omit $\boldsymbol{\theta}$ as an argument of I and write $I(p)$.

For obtaining the joint capacities for various models, we will extensively work with the binary entropy function. Again, this function can be quite ugly for arbitrary arguments α, but in some cases we can somewhat simplify the expressions. For instance, for arguments close to 0 or $\frac{1}{2}$ we have

$$h(\gamma) = \frac{\gamma(1-\ln\gamma)}{\ln 2} - O(\gamma^2) = O(\gamma \ln \gamma), \tag{94}$$

$$h\left(\frac{1}{2} \pm \gamma\right) = 1 - \frac{2\gamma^2}{\ln 2} - O(\gamma^4) = 1 - O(\gamma^2). \tag{95}$$

The most important properties to keep in mind are that $h(0) = h(1) = 0$ and h takes its maximum at $\alpha = \frac{1}{2}$ with $h(\frac{1}{2}) = 1$. Using only these latter properties, we immediately get the following lemma regarding deterministic attacks, i.e., attacks satisfying $\boldsymbol{\theta} \in \{0,1\}^{c+1}$.

Lemma 1: For any deterministic attack $\boldsymbol{\theta}$ satisfying the marking assumption $\theta_0 = 0$ and $\theta_c = 1$, the joint capacity equals $C^j(\boldsymbol{\theta}) = \frac{1}{c}$, and p is a maximizing value if it satisfies

$$a = \sum_{z:\ \theta_z = 1} \binom{c}{z} p^z (1-p)^{c-z} = \frac{1}{2}. \tag{96}$$

Proof: Since $\theta_z \in \{0,1\}$ for all z, we have $h(\theta_z) = 0$ for each z, so $a_h = 0$ and it thus follows that

$$C^j(\boldsymbol{\theta}) = \max_p \frac{1}{c}\left[h(a) - a_h\right] = \frac{1}{c}\max_p h(a). \tag{97}$$

Since $a = a(p)$ is continuous in p, and $a(0) = 0$ and $a(1) = 1$ due to the marking assumption, from the intermediate value theorem it follows that there must be a value $p \in (0,1)$ for which $a(p) = \frac{1}{2}$. So we get

$$C^j(\boldsymbol{\theta}) = \frac{1}{c}\max_p\left[h(a)\right] = \frac{1}{c}h\left(\frac{1}{2}\right) = \frac{1}{c}, \tag{98}$$

and p is a maximizing value iff $a(p) = \frac{1}{2}$. ■

This lemma makes finding the joint capacities and the optimal values of p very easy for several of the following models.

A. Fingerprinting

1) Interleaving attack: We previously saw that the simple capacity for the interleaving attack is proportional to $\frac{1}{c^2}$. The exact asymptotics for the joint capacity were previously derived by Huang and Moulin as well, showing that for large c the joint capacity is the same as the simple capacity.

Proposition 9: [32, Corollary 6] The joint capacity for the interleaving attack $\boldsymbol{\theta}_{\text{int}}$ is:

$$C^j(\boldsymbol{\theta}_{\text{int}}) = \frac{1}{2c^2 \ln 2} + O\left(\frac{1}{c^4}\right) \approx \frac{0.72}{c^2}. \tag{99}$$

Asymptotically, the interleaving attack is an "equalizing strategy" [32], guaranteeing that up to order terms $I(p)$ is the same for all $p \in (0,1)$.

2) All-1 attack: Since the all-1 attack is a deterministic attack and satisfies the marking assumption, the capacity follows immediately from Lemma 1, and finding the optimal value of p is straightforward.

Proposition 10: For the all-1 attack, the joint capacity and the maximizing value of p are:

$$C^j(\boldsymbol{\theta}_{\text{all1}}) = \frac{1}{c}, \qquad p^j_{\text{all1}} = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right). \tag{100}$$

To be precise, the optimal value for p is $p = 1 - 2^{-1/c}$.

3) Majority voting: Lemma 1 also applies to the majority voting attack $\boldsymbol{\theta}_{\text{maj}}$, and since this attack is symbol-symmetric the optimal value for p is trivially $p = \frac{1}{2}$.

Proposition 11: For the majority voting attack, the joint capacity and the corresponding optimal value of p are:

$$C^j(\boldsymbol{\theta}_{\text{maj}}) = \frac{1}{c}, \qquad p^j_{\text{maj}} = \frac{1}{2}. \tag{101}$$

Note that the joint capacity for majority voting is equal to the joint capacity of the all-1 attack, while their simple capacities differ. Also note that again, the optimal value for p is asymptotically the same as for the simple capacity.

4) Minority voting: Since minority voting is also a deterministic and symbol-symmetric attack, the following result directly follows from Lemma 1.

Proposition 12: The joint capacity and a corresponding optimal value of p for the minority voting attack are:

$$C^j(\boldsymbol{\theta}_{\text{min}}) = \frac{1}{c}\,, \qquad p^j_{\text{min}} = \frac{1}{2}\,. \tag{102}$$

In fact, there are three values of p that are asymptotically optimal, the other two being $p \approx \frac{\ln 2}{c}$ and $p \approx 1 - \frac{\ln 2}{c}$.

5) Coin-flip attack: Besides the interleaving attack, the only other non-trivial fingerprinting attack with respect to joint capacities is the coin-flip attack. This attack is not deterministic, so $a_h > 0$. Working out the details, we obtain the following result.

Proposition 13: For the coin-flip attack, the joint capacity and the maximizing value of p are:

$$C^j(\boldsymbol{\theta}_{\text{coin}}) = \frac{\log_2(5/4)}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.32}{c}\,, \tag{103}$$

$$p^j_{\text{coin}} = \frac{\ln(5/3)}{c} + O\left(\frac{1}{c^2}\right) \approx \frac{0.51}{c}\,. \tag{104}$$

Proof: For a_h, note that $h(\theta_0) = h(\theta_c) = 0$ and $h(\theta_z) = 1$ otherwise, so $a_h = 1 - p^c - (1-p)^c$. For a, recall from the proof of Proposition 5 that $a = \frac{1}{2}(1 - (1-p)^c + p^c)$. Combining the above, we get

$$I(p) = \frac{1}{c}\left[h\left(\frac{1-(1-p)^c+p^c}{2}\right) - (1 - p^c - (1-p)^c)\right]. \tag{105}$$

Since the attack is symbol-symmetric, w.l.o.g. we may assume that $p \leq \frac{1}{2}$, in which case the terms p^c are negligible for large c. Writing $t = 1 - (1-p)^c$, we get

$$I(t) = \frac{1}{c}\left[h\left(\frac{t}{2}\right) - t\right] + O\left(\frac{1}{c^2}\right). \tag{106}$$

This function has a maximum at $t = 1 - (1-p)^c = \frac{2}{5}$, which leads to the given values of p^j_{coin} and $C^j(\boldsymbol{\theta}_{\text{coin}})$. ∎

6) Unknown attacks: Finally, the case where the attack is not known (but is assumed to satisfy the marking assumption) was previously solved by Huang and Moulin, showing that again the interleaving attack is asymptotically optimal.

Proposition 14: [32, Corollary 7] The joint uninformed capacity is given by

$$C^j(\mathcal{P}_{\text{mark}}) = \frac{1}{2c^2 \ln 2} + O\left(\frac{1}{c^3}\right) \approx \frac{0.72}{c^2}\,, \tag{107}$$

and the optimizing encoder f_P and collusion channel $\boldsymbol{\theta}$ achieving this bound for large c are the arcsine distribution f_P^* and the interleaving attack $\boldsymbol{\theta}_{\text{int}}$.

So while the joint capacities are asymptotically the same as the simple capacities for the interleaving attack and for the uninformed setting, for several other attacks the joint capacities are strictly higher than the simple capacities.

B. Group testing

1) Classical model: Since the classical model is equivalent to the all-1 attack in group testing, the following result is immediate.

Corollary 2: For the classical group testing model, the joint capacity and the optimal value of p are:

$$C^j(\boldsymbol{\theta}_{\text{all1}}) = \frac{1}{c}\,, \qquad p^j_{\text{all1}} = \frac{\ln 2}{c} + O\left(\frac{1}{c^2}\right). \tag{108}$$

This result was previously derived by Sebő [48, Theorem 2], who also showed that $p = 1 - 2^{-1/c} \approx \frac{\ln 2}{c}$ is optimal.

2) Additive noise: The additive noise model described in Section II-B2 was previously studied in the context of capacities in e.g. [7], [18], [49]. Cheraghchi et al. [18] showed that $C^j(\boldsymbol{\theta}_{\text{add}}) = O(\frac{(1-r)^3}{c})$, while Atia and Saligrama [7] showed that $C^j(\boldsymbol{\theta}_{\text{add}}) = O(\frac{1-r}{c})$. Looking closely at their proof, they show that one obtains a capacity of $I(p) \geq \frac{1-r}{ec\ln 2} \approx \frac{1.88(1-r)}{c}$ using $p = \frac{1}{c}$ for large c.[4]

Below we improve upon these results, by (i) providing the exact leading constant on the capacity; (ii) showing exactly how the first order term (in r) scales for small r; and (iii) showing how p scales in terms of r.

Proposition 15: For the additive noise model, the joint capacity and the corresponding optimal value of p are:

$$C^j(\boldsymbol{\theta}_{\text{add}}) = \frac{1}{c}\left(1 - \tfrac{1}{2}h(r) + O(r^2)\right) + O\left(\frac{1}{c^2}\right), \tag{109}$$

$$p^j_{\text{add}} = \frac{\ln 2}{c}\left(1 - \frac{r(1+\ln r)}{2\ln 2} + O(r^2)\right) + O\left(\frac{1}{c^2}\right). \tag{110}$$

Proof: First, from the definition of $\boldsymbol{\theta}_{\text{add}}$ it follows that $a = 1 - (1-p)^c(1-r)$, $h(\theta_0) = h(1-r)$ and $h(\theta_z) = 0$ for $z > 0$. So the mutual information satisfies

$$I(p) = \frac{1}{c}\left[h((1-p)^c(1-r)) - (1-p)^c h(1-r)\right]. \tag{111}$$

Writing $s = (1-p)^c(1-r)$ this can be simplified to

$$I(s) = \frac{1}{c}\left[h(s) - \frac{h(r)s}{1-r}\right]. \tag{112}$$

We want to maximize I, so we take the derivative with respect to s, and set it equal to 0 to obtain a condition for s, and hence for p:

$$\log_2\left(\frac{s}{1-s}\right) = -\frac{h(r)}{1-r}\,. \tag{113}$$

For small r, the right hand side goes to 0, which implies that s is close to $\frac{1}{2}$. So assuming r is small, we let $s = \frac{1}{2}(1+\varepsilon)$ and obtain the following Taylor expansion for the left hand side:

$$\log_2\left(\frac{1+\varepsilon}{1-\varepsilon}\right) = \log_2\left(1 + 2\varepsilon + O(\varepsilon^2)\right) = \frac{2\varepsilon}{\ln 2} + O(\varepsilon^2). \tag{114}$$

[4]The authors of [7] confirmed that the formula below [7, (45)] contains a typo: there should be an extra e in the numerator of the code length T.

This means that for small r, the optimal choice for ε is

$$\varepsilon = -\frac{h(r)\ln 2}{2} + O(r^2). \tag{115}$$

So for s we obtain

$$s = (1-p)^c(1-r) = \frac{1}{2}\left(1 - \tfrac{1}{2}h(r)\ln 2 + O(r^2)\right). \tag{116}$$

Substituting s into $I(s)$, and solving s for p, we obtain the given values for p^j_{add} and $C^j(\boldsymbol{\theta}_{\text{add}})$. ∎

Note that this means that any valid group testing algorithm asymptotically requires at least the following number of tests:

$$\ell \geq \frac{c\log_2 n}{1 - \frac{1}{2}h(r) + O(r^2)}\left(1 + O\left(\frac{1}{c}\right)\right). \tag{117}$$

Since $r = o(h(r))$ for small r, this shows that the result of [7] is slightly off; due to their suboptimal choice of p, they obtained a code length which scales "better" in r, but has a higher leading constant and thus converges to the wrong limit.

3) Dilution noise: The dilution noise model, as described in Section II-B3, was previously studied in the context of lower bounds by Atia and Saligrama [7]. In terms of capacities, they showed that for large c, one has $C^j(\boldsymbol{\theta}_{\text{dil}}) = O(\frac{(1-r)^2}{c})$. Again, they were not interested in leading constants, so they fixed p to the suboptimal choice $p = \frac{1}{c}$. We improve upon their result by finding the leading constant explicitly, and proving how p^j_{dil} and $C^j(\boldsymbol{\theta}_{\text{dil}})$ scale in terms of r.

Proposition 16: For the dilution noise model with parameter r, the joint capacity and the corresponding maximizing value of p are:

$$C^j(\boldsymbol{\theta}_{\text{dil}}) = \frac{1}{c}\left(1 - \frac{\ln 2}{2}h(r) + O(r^2)\right) + O\left(\frac{1}{c^2}\right), \tag{118}$$

$$p^j_{\text{dil}} = \frac{\ln 2}{c}\left(1 + r - \frac{1-\ln 2}{2}h(r) + O(r^2)\right) + O\left(\frac{1}{c^2}\right) \tag{119}$$

Proof: For this attack, we have $\theta_z = 1 - r^z$. Let us first look at $h(a)$:

$$h(a) = h\left(\sum_{z=0}^{c}\binom{c}{z}p^z(1-p)^{c-z}(1-r^z)\right) \tag{120}$$

$$= h(1 - (1-p+pr)^c). \tag{121}$$

Next, consider a_h:

$$a_h = \sum_{z=1}^{c}\binom{c}{z}p^z(1-p)^{c-z}h(1-r^z). \tag{122}$$

For small r, the only significant contribution to the sum comes from the term with $z = 1$:

$$a_h = cp(1-p)^{c-1}h(r) + O(r^2). \tag{123}$$

The optimal value of p again lies close to $\frac{\ln 2}{c}$; in particular, the value is mostly determined by the term $h((1-p+pr)^c)$, which has a maximum at $(1-p+pr)^c = \frac{1}{2}$. Writing $(1-p+pr)^c = \frac{1}{2}(1+\varepsilon)$, we have

$$p = \frac{1}{c}\left(\ln 2 + r\ln 2 - \varepsilon - r\varepsilon + \frac{\varepsilon^2}{2} + O(r^2, \varepsilon^2 r, \varepsilon^3)\right), \tag{124}$$

$$(1-p)^c = \frac{1}{2}\left(1 - r\ln 2 + \varepsilon + r\varepsilon - \frac{\varepsilon^2}{2} + O(r^2, \varepsilon^2 r, \varepsilon^3)\right). \tag{125}$$

This means that $I(p) = I(\varepsilon)$ satisfies (neglecting terms of the order $r^2, \varepsilon^2 r, \varepsilon^3, c^{-1}$)

$$I(\varepsilon) \sim 1 - \frac{1}{2}h(r)\ln 2 + \frac{1}{2}\varepsilon h(r)(1-\ln 2) - \frac{\varepsilon^2}{2\ln 2}. \tag{126}$$

Taking the derivative with respect to ε and setting it equal to 0, we obtain

$$\varepsilon = \frac{1}{2}h(r)\ln 2(1-\ln 2) + O(r^2). \tag{127}$$

Substituting this value for ε in the expressions for p and I, we get the results. ∎

For the resulting lower bound on the code length ℓ, one thus obtains

$$\ell \sim \frac{c\log_2 n}{1 - \frac{1}{2}h(r)\ln 2 + O(r^2)}. \tag{128}$$

So also in the dilution noise model, the first order term in the denominator scales as $h(r)$ rather than r, as one might suspect from the results of [7].

4) Threshold without gaps: For threshold group testing with $u = l+1$ (as described in Section II-B4) we now consider two different cases for u: $u = \Theta(c)$ and $u = o(c)$. In both cases, the capacity follows directly from Lemma 1, but we can obtain slightly more accurate asymptotics for p in the second case. The first case is sometimes referred to in the literature as majority group testing [1]–[3].

Proposition 17: For the threshold group testing model with $u = \ell + 1$, the joint capacity is $\frac{1}{c}$, and the corresponding maximizing value of p is:

$$u = \Theta(c): \quad p^j_{\text{thr}}[\boldsymbol{\theta}^{(u)}_{\text{thr}}] = \frac{1}{c}(u + \gamma) \qquad (|\gamma| \leq 1) \tag{129}$$

$$u = o(c): \quad p^j_{\text{thr}}[\boldsymbol{\theta}^{(u)}_{\text{thr}}] = \frac{1}{c}\left(u - \frac{1}{3} + O\left(\frac{1}{u}\right)\right). \tag{130}$$

Proof: From Lemma 1 it follows that the capacity is $\frac{1}{c}$ and that the optimal value of p satisfies $a = \frac{1}{2}$. Writing out a, we have

$$a = \sum_{z=0}^{u-1}\binom{c}{z}p^z(1-p)^{c-z} = \frac{1}{2}. \tag{131}$$

The fact that $a = \frac{1}{2}$ roughly means that u is the median of the binomial distribution with c trials and probability of success p. Since the median of a binomial distribution is one of the two integers closest to cp, it follows that $|u - cp| \leq 1$ leading to the result for the case $u = \Theta(c)$.

For the case $u = o(c)$, note that $p = O(\frac{1}{c})$, so $(1-p)^z = 1 - O(p)$ for $z < u$. So we can expand a around $c = \infty$ as:

$$a = (1-p)^c\sum_{z=0}^{u-1}\binom{c}{z}p^z + O\left(\frac{1}{c}\right). \tag{132}$$

Since the solution is in the range $p = \Theta(\frac{1}{c})$, let us write $p = \frac{\alpha}{c}$ for some constant α. A Taylor expansion around $c = \infty$ of the binomial coefficients then gives us

$$a = e^{-\alpha} \sum_{z=0}^{u-1} \frac{\alpha^z}{z!} + O\left(\frac{1}{c}\right). \tag{133}$$

The condition that $a = \frac{1}{2}$ means that asymptotically, $u - 1$ is the median of the Poisson distribution with parameter $\lambda = \alpha$. Using results about the median of the Poisson distribution [20], we obtain

$$\alpha = u - \frac{1}{3} + O\left(\frac{1}{u}\right). \tag{134}$$

Substituting this back into p, we get the result. ∎

Note that for $u = 1$ and $c \to \infty$, the above approximation says $p \approx \frac{0.67}{c}$, when in reality the optimum is at $p \sim \frac{\ln 2}{c} \approx \frac{0.69}{c}$, showing that already for small values of u the term $u - \frac{1}{3}$ is quite accurate.

5) Threshold with gaps: For threshold group testing with gaps, let us again consider the two models described in Section II-B5: the coin-flip gap model and the interleaving gap model. For both models, we can again interpolate between results obtained earlier in this section to obtain estimates for $C^j(\boldsymbol{\theta}^{(l,u)}_{\text{coin}})$ and $C^j(\boldsymbol{\theta}^{(l,u)}_{\text{int}})$ for various l and u, and verify our intuition numerically (see Figure 2). In both plots, from Proposition 17 it follows that the diagonals have value $c \cdot C^j(\boldsymbol{\theta}) = 1$, while the upper left corner in Figure 2a converges to $\log_2(5/4) \approx 0.32$ (Proposition 13) and the upper left corner of Figure 2b converges to 0 (Proposition 9). In the left graph, even for small gaps we see that the capacity quickly decreases and approaches the coin-flip capacity. In the right graph, we see that the capacity decreases more gradually as the gap size increases.

IV. Discussion

Building upon previous work of Huang and Moulin and working our way through the resulting expressions for the capacities, we have derived explicit asymptotics for both the simple and joint capacities for various fingerprinting and group testing models. In the end the results from fingerprinting turned out to be useful in threshold group testing as well, for understanding the numerics of Figures 1 and 2 and estimating the capacities for various threshold group testing models.

One important result with respect to group testing is that the simple capacity in the traditional model is asymptotically a factor $\log_2(e)$ lower than the joint capacity. While the joint capacity was well known, to the best of our knowledge the simple capacity had not yet been derived before. This result shows that efficient (simple) group testing algorithms will never be able to achieve the code lengths of optimal joint decoders, and that various existing methods (e.g. [12], [13]) are suboptimal, even for simple decoding. The related paper [40] explicitly shows how the bounds on the code lengths of simple decoders can be attained with log-likelihood decoders.

Comparing the simple and joint capacities, another result worth mentioning is that except for in the cases previously analyzed by Huang and Moulin, there is always a gap between the simple and joint capacities. In fingerprinting, this means that if the pirates use a suboptimal attack, joint decoders are asymptotically significantly better than simple decoders. In terms of group testing, this means that in almost all models, simple decoders are strictly worse than joint decoders. So although joint decoders are generally slower, the benefits of joint decoding (a much shorter code length) may outweigh the costs of a higher decoding complexity.

V. Open problems

Let us finish by mentioning some open problems which are left for future work.

A. Dynamic fingerprinting and adaptive group testing

While this paper considered only static fingerprinting and non-adaptive group testing, in some settings the feedback Y may be obtained in real-time. For instance, in pay-tv pirates may try to duplicate a fingerprinted broadcast, while in group testing it may sometimes be possible to do group tests sequentially. These dynamic or adaptive games have received considerable attention as well [4], [8], [23], [27], [35], [36], [38] but little is known about the capacities of these games. Are the dynamic/adaptive capacities strictly higher than the static/non-adaptive capacities in the probabilistic model considered in this paper?

B. Tuple decoders and tuple capacities

Recall that simple decoders base their decisions only on individual code words, while joint decoders base their decisions on the entire code $\mathcal{X}$. The extra information used by joint decoders generally causes the joint capacity to be higher than the simple capacity, but the complexity of decoding may be higher as well. A possible way to obtain a trade-off between the code length and the time complexity would be *tuple decoding*: basing the decision to accuse a user j only on tuples of size at most t. This could be seen as a generalization of simple and joint decoding, since those models correspond to $t = 1$ and $t = c$ respectively. Such decoders were previously considered in e.g. [5], [43], [44], [46], and an obvious question is: can we somehow quantize this trade-off between the time complexity and the code length? And can we formally derive capacities for this tuple decoding model?

C. Non-binary codes in fingerprinting

A common generalization in fingerprinting is to assume that symbols come from an alphabet of size $q \geq 2$, rather than assuming that the code $\mathcal{X}$ is a binary code ($q = 2$). This generalization was considered in e.g. [10], [33], [45], [46], [50]. In the uninformed fingerprinting game, the capacity decreases linearly with q [10], [33], so there may be significant benefits going from a binary to a q-ary alphabet. For the models considered in this paper, for which the capacity is only linear in c, it is easy to see that the capacity cannot increase linearly with q. Some basic numerics seem to indicate that the capacity increases with a factor $\log q$, but a more detailed analysis is required.

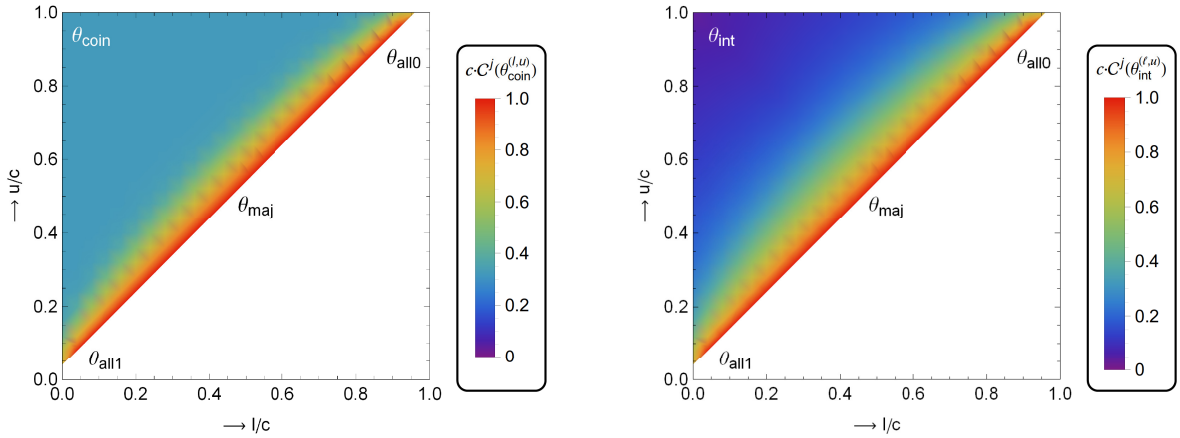

(a) Joint capacity in the coin-flip gap model

(b) Joint capacity in the interleaving gap model

Fig. 2. The joint capacity for threshold group testing with different values of l and u, when there is (a) a coin-flip gap or (b) an interleaving gap. The three corners correspond to the all-1, all-0 and coin-flip (a) or interleaving (b) attack, and the center of the graphs corresponds to the majority voting attack in fingerprinting. The capacity is maximal on the diagonal, for which $c \cdot C^j(\boldsymbol{\theta}) = 1$.

D. Universal encoding in fingerprinting

Finally, instead of assuming that the pirate attack is known in advance, in fingerprinting it is more often assumed that the encoding is done for arbitrary attacks, and that only the decoding step may be tuned to fit the pirate attack [15], [28], [42], [45]. Since the asymptotically optimal universal encoding strategy is to use the arcsine distribution f_P^* for generating biases p, one could try deriving the capacities for the various fingerprinting attacks in case the distribution f_P is fixed in advance as $f_P \equiv f_P^*$. Previous results [45] showed that the capacities probably scale as $c^{-3/2}$, and numerics of the associated capacities (Figure 3) seem to verify this. Obtaining exact expressions for the simple and joint capacities under 'universal encoding' is left for future work.

ACKNOWLEDGMENTS

The author is grateful to Benne de Weger for his help with some of the proofs in this paper, and for his comments on drafts of this manuscript that helped improve the paper. The author would further like to thank Jeroen Doumen, Teddy Furon, Jan-Jaap Oosterwijk, and Boris Škorić for their valuable comments and suggestions.

REFERENCES

[1] R. Ahlswede, C. Deppe, and V. S. Lebedev, "Bounds for Threshold and Majority Group Testing," *IEEE International Symposium on Information Theory (ISIT)*, pp. 69–73, 2011.

[2] R. Ahlswede, C. Deppe, and V. S. Lebedev, "Majority Group Testing with Density Tests," *IEEE International Symposium on Information Theory (ISIT)*, pp. 326–330, 2011.

[3] R. Ahlswede, C. Deppe, and V. S. Lebedev, "Threshold and Majority Group Testing," *Information Theory, Combinatorics, and Search Theory*, LNCS vol. 7777, pp. 488–508, 2013.

[4] M. Aldridge, "Adaptive Group Testing as Channel Coding with Feedback," *IEEE International Symposium on Information Theory (ISIT)*, pp. 1832–1836, 2012.

[5] E. Amiri and G. Tardos, "High Rate Fingerprinting Codes and the Fingerprinting Capacity," *20th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 336–345, 2009.

[6] G. K. Atia and V. Saligrama, "Noisy Group Testing: An Information Theoretic Perspective," *47th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 355–362, 2009.

[7] G. K. Atia and V. Saligrama, "Boolean Compressed Sensing and Noisy Group Testing," *IEEE Transactions on Information Theory*, vol. 58, no. 3, pp. 1880–1901, 2012.

[8] L. Baldassini, O. Johnson, and M. Aldridge, "The Capacity of Adaptive Group Testing," *arXiv*, 2013.

[9] W. Berchtold and M. Schäfer, "Performance and Code Length Optimization of Joint Decoding Tardos Fingerprinting," *ACM Symposium on Multimedia and Security (MMSec)*, pp. 27–32, 2012.

[10] D. Boesten and B. Škorić, "Asymptotic Fingerprinting Capacity for Non-Binary Alphabets," *13th Conference on Information Hiding (IH)*, pp. 1–13, 2011.

[11] D. Boneh and J. Shaw, "Collusion-Secure Fingerprinting for Digital Data," *IEEE Transactions on Information Theory*, vol. 44, no. 5, pp. 1897–1905, 1998.

[12] C.-L. Chan, P. H. Che, S. Jaggi, and V. Saligrama, "Non-adaptive probabilistic group testing with noisy measurements: Near-optimal bounds with efficient algorithms," *49th Allerton Conference on Communication, Control, and Computing*, pp. 1832–1839, 2011.

[13] C.-L. Chan, S. Jaggi, V. Saligrama, and S. Agnihotri, "Non-Adaptive Group Testing: Explicit Bounds and Novel Algorithms," *IEEE International Symposium on Information Theory (ISIT)*, pp. 1837–1841, 2012.

[14] C.-L. Chan, S. Cai, M. Bakshi, S. Jaggi, and V. Saligrama, "Near-Optimal Stochastic Threshold Group Testing," *arXiv*, 2013.

[15] A. Charpentier, F. Xie, C. Fontaine, and T. Furon, "Expectation Maximization Decoding of Tardos Probabilistic Fingerprinting Code," *SPIE Proceedings / Media Forensics and Security*, vol. 7254, 2009.

[16] H.-B. Chen and H.-L. Fu, "Nonadaptive Algorithms for Threshold Group Testing," *Discrete Applied Mathematics*, vol. 157, no. 7, pp. 1581–1585, 2009.

[17] M. Cheraghchi, A. Hormati, A. Karbasi, and M. Vetterli, "Compressed Sensing with Probabilistic Measurements: A Group Testing Solution," *47th Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 30–35, 2009.

[18] M. Cheraghchi, A. Hormati, A. Karbasi, and M. Vetterli, "Group Testing with Probabilistic Tests: Theory, Design and Application," *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 7057–7067, 2011.

[19] M. Cheraghchi, "Improved Constructions for Non-Adaptive Threshold Group Testing," *Algorithmica*, vol. 67, no. 3, pp. 384–417, 2013.

[20] K. P. Choi, "On the Medians of Gamma Distributions and an Equation of Ramanujan," *Proceedings of the American Mathematical Society*, vol. 121, no. 1, pp. 245–251, 1994.

[21] T. M. Cover and J. A. Thomas, *Elements of Information Theory (2nd Edition)*, Wiley Press, 2006.

[22] P. Damaschke, "Threshold Group Testing," *General Theory of Information Transfer and Combinatorics*, LNCS vol. 4123, pp. 707–718, 2006.

[23] A. De Bonis, L. Gasieniec, and U. Vaccaro, "Optimal Two-Stage Algorithms for Group Testing Problems," *SIAM Journal on Computing*, vol. 34, no. 5, pp. 1253–1270, 2005.

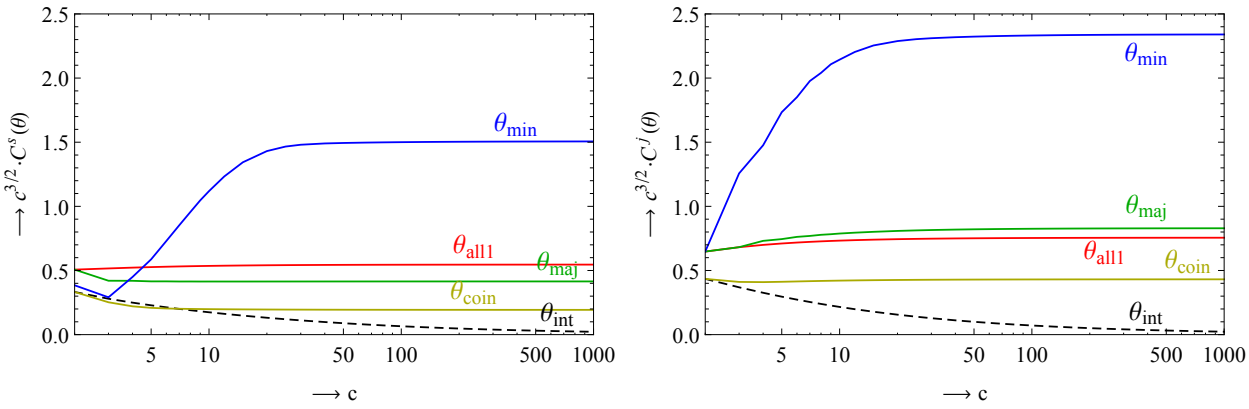

(a) Simple capacity in the universal encoding model

(b) Joint capacity in the universal encoding model

Fig. 3. The simple (a) and joint (b) universal encoding capacities (multiplied by $c^{3/2}$) as a function of c, for different pirate attacks. Except for the interleaving attack, for which the capacity scales as c^{-2} (the dashed line), these capacities all seem to scale as $c^{-3/2}$.

[24] R. Dorfman, "The Detection of Defective Members of Large Populations," *The Annals of Mathematical Statistics*, vol. 14, no. 4, pp. 436–440, 1943.

[25] A. G. D'yachkov and V. V. Rykov, "Bounds on the length of disjunctive codes," *Problemy Peredachi Informatsii*, vol. 18, no. 3, pp. 7–13, 1982.

[26] A. G. D'yachkov, V. V. Rykov, and A. M. Rashad, "Superimposed distance codes," *Problems of Control and Information Theory*, vol. 18, no. 4, pp. 237–250, 1989.

[27] A. Fiat and T. Tassa, "Dynamic Traitor Tracing," *Journal of Cryptology*, vol. 14, no. 3, pp. 211–223, 2001.

[28] T. Furon and L. Pérez-Freire, "EM Decoding of Tardos Traitor Tracing Codes," *ACM Symposium on Multimedia and Security (MM&Sec)*, pp. 99–106, 2009.

[29] Y.-W. Huang and P. Moulin, "Capacity-Achieving Fingerprint Decoding," *IEEE Workshop on Information Forensics and Security (WIFS)*, pp. 51–55, 2009.

[30] Y.-W. Huang and P. Moulin, "Saddle-Point Solution of the Fingerprinting Capacity Game under the Marking Assumption," *IEEE International Symposium on Information Theory (ISIT)*, pp. 2256–2260, 2009.

[31] Y.-W. Huang and P. Moulin, "Maximin Optimality of the Arcsine Fingerprinting Distribution and the Interleaving Attack for Large Coalitions," *IEEE Workshop on Information Forensics and Security (WIFS)*, pp. 1–6, 2010.

[32] Y.-W. Huang and P. Moulin, "On the Saddle-Point Solution and the Large-Coalition Asymptotics of Fingerprinting Games," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 1, pp. 160–175, 2012.

[33] Y.-W. Huang and P. Moulin, "On Fingerprinting Capacity Games for Arbitrary Alphabets and Their Asymptotics," *IEEE International Symposium on Information Theory (ISIT)*, pp. 2571–2575, 2012.

[34] F. K. Hwang, "Group Testing with a Dilution Effect," *Biometrika*, vol. 63, no. 3, pp. 671–680, 1976.

[35] T. Laarhoven, J.-J. Oosterwijk, and J. Doumen, "Dynamic Traitor Tracing for Arbitrary Alphabets: Divide and Conquer," *IEEE Workshop on Information Forensics and Security (WIFS)*, pp. 240–245, 2012.

[36] T. Laarhoven, J. Doumen, P. Roelse, B. Škorić, and B. de Weger, "Dynamic Tardos Traitor Tracing Schemes," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4230–4242, 2013.

[37] T. Laarhoven, "Efficient Probabilistic Group Testing Based on Traitor Tracing," *51st Annual Allerton Conference on Communication, Control and Computing (Allerton)*, pp. 1458–1465, 2013.

[38] T. Laarhoven, "Dynamic Traitor Tracing Schemes, Revisited," *IEEE Workshop on Information Forensics and Security (WIFS)*, pp. 191–196, 2013.

[39] T. Laarhoven, "Capacities and Capacity-Achieving Decoders for Various Fingerprinting Games," *ACM Workshop on Information Hiding and Multimedia Security (IH&MMSec)*, 2014. This is a preliminary version of the present paper.

[40] T. Laarhoven, "Asymptotics of Fingerprinting and Group Testing: Capacity-Achieving Log-Likelihood Decoders," *submitted to IEEE Transactions on Information Theory*, 2014.

[41] V. S. Lebedev, "Separating Codes and a New Combinatorial Search Model," *Problems of Information Transmission*, vol. 46, no. 1, pp. 1–6, 2010.

[42] P. Meerwald and T. Furon, "Group Testing Meets Traitor Tracing," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4204–4207, 2011.

[43] P. Meerwald and T. Furon, "Toward Practical Joint Decoding of Binary Tardos Fingerprinting Codes," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 4, pp. 1168–1180, 2012.

[44] P. Moulin, "Universal Fingerprinting: Capacity and Random-Coding Exponents," *arXiv:0801.3837v3 [cs.IT]*, 2011.

[45] J.-J. Oosterwijk, B. Škorić, and J. Doumen, "A Capacity-Achieving Simple Decoder for Bias-Based Traitor Tracing Schemes," *Cryptology ePrint Archive*, 2013.

[46] J.-J. Oosterwijk, J. Doumen, and T. Laarhoven, "Tuple Decoders for Traitor Tracing Schemes," *SPIE Proceedings*, vol. 9028, 2014.

[47] L. Pérez-Freire and T. Furon, "Blind Decoder for Binary Probabilistic Traitor Tracing Codes," *IEEE Workshop on Information Forensics and Security (WIFS)*, pp. 46–50, 2009.

[48] A. Sebő, "On Two Random Search Problems," *Journal of Statistical Planning and Inference*, vol. 11, pp. 23–31, 1985.

[49] D. Sejdinovic and O. Johnson, "Note on Noisy Group Testing: Asymptotic Bounds and Belief Propagation Reconstruction," *48th Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 998–1003, 2010.

[50] B. Škorić, S. Katzenbeisser, and M. U. Celik, "Symmetric Tardos Fingerprinting Codes for Arbitrary Alphabet Sizes," *Designs, Codes and Cryptography*, vol. 46, no. 2, pp. 137–166, 2008.

[51] D. R. Stinson, T. van Trung, and R. Wei, "Secure Frameproof Codes, Key Distribution Patterns, Group Testing Algorithms and Related Structures," *Journal of Statistical Planning and Inference*, vol. 86, no. 2, pp. 595–617, 2000.

[52] G. Tardos, "Optimal Probabilistic Fingerprint Codes," *35th ACM Symposium on Theory of Computing (STOC)*, pp. 116–125, 2003.